

**HIGH SPEED INFORMATION PROCESSING AND MASS  
STORAGE SYSTEM AND METHOD, PARTICULARLY FOR  
INFORMATION AND APPLICATION SERVERS**

**Related Applications:**

This application is related to Patent Application No. PCT/US99/05231 of Richard Dellacona filed March 10, 1999, which is based upon Provisional Patent Application Serial No. 60/077,643, filed March 10, 1998, and to U.S. Patent Application Serial No. 09/071282 of Richard Dellacona filed May 1, 1998, all of which are hereby incorporated herein by reference.

**BACKGROUND OF THE INVENTION**

**Field of the Invention**

This invention relates to a high speed, microcomputer based, Fibre Channel compatible and fault tolerant information processing and mass storage system especially suited for information servers and application servers. In particular, the present invention relates to a method and apparatus for information processing and storage involving a unique and extremely versatile system architecture, including a dual loop arbitrated, Fibre Channel capable, multiple-fault tolerant, hot-swappable mass storage disk array and including a method and apparatus for providing an enterprise-wide information or application server system using such disk array.

**State of the Art**

Efforts have been made in the past to provide a mass storage file server capable of delivering information throughout an enterprise with high speed data throughput, scalable data storage capability in a convenient, easily configurable enclosure using well known, industry standard operating software and hardware. However, such systems have typically experienced many shortcomings and problems associated with equipment incompatibility as well as with the inability of presently available computer and communications hardware to sustain performance and service failure of component devices. Such shortcomings have included the lack of capability to add storage devices to accommodate increased storage requirements

or to replace failed storage devices without the need to completely power down the information server. Some of the compatibility problems have involved, for example, bottlenecks in sharing information among the equipment components of various vendors. The above-referenced Dellacona patent applications address some of these problems and others, and provide unique solutions which are described and claimed therein.

The present invention further addresses some of the problems discussed in the referenced Dellacona applications, as well as others. For example, the present invention further addresses the problem of scalability and customization in information processing and storage systems for different applications. In some applications, there may be a greater need for processing capability rather than storage capacity, while other applications may require just the opposite. Yet other applications may require storage expansion for existing information processing systems. This invention provides an architecture which will readily accommodate such needs.

In addition, mass storage systems can create considerable heat, particularly where they are disk drive based. If the heat is not effectively removed, it can affect the reliability and life of the system. Often, it is difficult to remove the heat because of obstructions caused by the physical configuration of back planes and mid planes which act as barriers to air flow. Typically, for example, all of the disk drives of a mass storage module or array are typically plugged into connectors on the face of a backplane or mid plane that extends across the entire module. Whether enclosed in a cabinet or other enclosure or rack mounted without an enclosure, air flow through the module is inhibited by this type of structural arrangement, and excessive heating can occur, particularly in the vicinity of the disk drives.

Also, it is desirable to be able to hot swap individual disk drives of a mass storage module to accommodate the need for more storage capacity, but the system storage requirements may outgrow the capacity of the module and it may also be desirable to have the capability of adding modules without powering down the system. Of course, this capability must be provided without disturbing the operation of the existing module and with a minimum of signal degradation as modules are added.

### SUMMARY OF THE INVENTION

The present invention obviates one or more of the foregoing problems and/or shortcomings of the prior art through the provision of an information processing and mass storage method and system, including a unique mass storage array, particularly  
5 suited for information servers or application servers, with a novel system architecture which permits the addition or replacement of storage devices without interrupting or seriously degrading the operation of the system and which is highly fault-tolerant and reliable. In addition, the invention obviates one or more of the foregoing problems by providing a novel physical layout that permits the effective  
10 removal of heat from a system module containing heat creating components.

In accordance with one embodiment of the invention, an information processing and mass storage system adapted to be readily expandable to increase its storage capacity while the system is in operation comprises at least one module containing (a) at least one computer, (b) a plurality of plug-in storage devices such  
15 as disk drives for storing information, (c) a storage device bypass circuit board associated with each storage device, with each storage device being plugged into a connector on the bypass circuit board, (d) a module bypass circuit board including an optical input/ output connector for outputting electrical signals from the module as light signals and for inputting light signals into the module as electrical signals  
20 and (e) a controller connecting the computer with each of the storage devices through its associated storage device bypass circuit board and through the module bypass circuit board.

Certain information server configurations in accordance with the invention include one or more computers, each computer connected to communicate through a  
25 Fibre Channel controller with a mass storage array comprising a plurality of bypass circuit boards, at least some of which are connected to an information storage device. In one embodiment, the controller provides a dual loop communications channel comprising two complete communication paths to each of bypass circuit boards and associated storage devices. In another embodiment, the controller  
30 provides a single loop which traverses the bypass circuit boards and any associated storage devices twice. In a one configuration with two or more computers, the Fibre

00279721051200

In a preferred embodiment the computer is preferably a suitable conventional single board computer. The controller preferably is a conventional arbitrated dual channel Fibre Channel system through which the computer communicates with each of the storage device bypass circuit boards and the module bypass circuit board. The bypass circuit boards may be any suitable circuits which form a continuous loop for the Fiber Channel controller regardless of whether a disk drive is plugged into the drive bypass circuit board. The loop continues through other modules when they are connected to the module bypass circuit thereby readily permitting expansion while maintaining a unitary information processing and mass storage system.

In accordance with yet another embodiment of the present invention, a high speed information processing and/or mass storage system with disk drives for information storage includes at least one module with a plurality of drive bypass circuit boards, each including a drive bypass circuit board connector. At least one opening is provided between connectors to permit the flow of air between the connectors. Each drive bypass circuit board is a relatively flat circuit board with connectors on different edges of the board, wherein one of the connectors is the connector which receives the disk drive and the other connector connects to said drive bypass circuit board connector. The connectors, bypass circuit boards and drives are arranged such that when they are connected there is a path for air flow from outside the module alongside each bypass circuit board and its associated disk

drive for cooling purposes without any backplane obstruction. Where the mass storage system is housed in an enclosure, at least one fan is mounted to force air from outside said enclosure through the spaces between said bypass boards and drives, preferably through the openings between the drive bypass circuit board connectors.

It will be appreciated that the present invention provides a novel high speed information processing and/or mass storage system particularly suitable for information and application servers. The system is scalable, fault tolerant, and reliable both because of its novel system architecture and its physical layout. Other features and advantages of the invention will become apparent from the following detailed description of exemplary and preferred embodiments when read in conjunction with the drawings which illustrate, by way of example, the principles of the invention.

#### **BRIEF DESCRIPTION OF THE DRAWINGS**

Figure 1 is a functional block diagram generally illustrating an information or application server system according to the present invention using a high speed mass storage system of the invention;

Figures 2A, 2B and 2C are diagrammatic representations illustrating various system configurations for different applications wherein different ratios of processing and storage are required;

Figure 3 is a functional block diagram illustrating a mass storage module of Figure 1 in greater detail;

Figure 4A is a functional block diagram illustrating the dual loop communication path from the Fibre Channel controller of Figure 3 connected so as to provide one logical Fibre Channel communication path traversing each of the storage device bypass circuit boards twice;

Figure 4B is a functional block diagram illustrating the dual loop communication path from the Fibre Channel controller of Figure 3 connected so as to provide two logical Fibre Channel communication paths traversing each of the storage device bypass circuit boards once, thereby being independently available to communicate with each of the storage devices;

Figure 5 is a functional block diagram illustrating an information or application server system according to the present invention wherein two or more servers are connected to a mass storage array which may include one or more mass storage modules;

5           Figure 6 is a functional block diagram illustrating an embodiment of a web server application configured in accordance with the architecture of the present invention;

          Figure 7 is a functional block diagram illustrating an embodiment of a basic video streaming application configured in accordance with the architecture of the present invention;

10           Figure 8 is a functional block diagram illustrating another embodiment of a video streaming application configured in accordance with the architecture of the present invention and including a duplicated index server;

          Figure 9 is a functional block diagram illustrating another embodiment of a video streaming application configured in accordance with the architecture of the present invention and including a distributed index server;

15           Figure 10 is a functional block diagram illustrating one embodiment of a storage device bypass circuit board according to the present invention;

          Figure 11 is a functional block diagram illustrating one embodiment of a chassis bypass circuit board according to the present invention;

20           Figures 12 A and 12 B are diagrammatic representations of the physical layout of a drive bypass circuit board and disk drive illustrating the preferred connector arrangement according to the present invention; and

          Figures 13 A and 13 B are diagrammatic representations of a connector arrangement for connecting a single board computer (SBC) to various input/out devices according to the present invention.

#### **DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS**

One embodiment of an information or application server system in accordance with the present invention using the high speed information processing and mass storage system of the invention is illustrated in Figure 1. Referring to 30   Figure 1, the server, designated generally by the reference numeral 100, includes a computer 102, a controller 104, preferably a Fibre Channel controller, and a

005027250500

communications interface or access card 106. The server 100 communicates via the Fibre Channel controller 104 with a mass storage array 200 which includes one or more mass storage modules 202A ... 202n. As illustrated, the computer 102 may also communicate with a suitable diagnostic computer 110 as is described in the  
5   aforementioned Dellacona applications.

          The computer 102 preferably comprises single board high speed computer running a computer industry standard operating system software program such as, for example, Windows NT available from Microsoft Corporation. An operating system like Windows NT may be stripped down to remove those elements of the  
10   program which are not needed, if desired to preserve memory or to increase operating speed. Suitable conventional drivers of the type used for similar applications may be provided as necessary to support the particular architecture being implemented.

          The computer may include a display such as a touch screen display, and  
15   various storage and peripheral devices (not shown) as required. The single board computer can include any of a wide number of suitable devices including, but not limited to, the Compact PCI CPU Board with Pentium Processor, Model No. ZT 5510, available from Ziatech Corporation. Modifications to enhance performance of the ZT 5510 can include an onboard 40 MB flash memory card for permanent  
20   storage of the non-reconfigurable portions of the Windows NT operating system software and an onboard, removable, PCMCIA 40 Mb flash memory card, "D2 FlashDisk" available from Sandisk Corporation for read/writeable storage of the reconfigurable portions of the Windows NT software.

          The Fibre Channel controller 104 may be any suitable design according to  
25   the Fibre Channel Consortium created as a separate board or incorporated into the single board computer design. The communications interface or access card 106 may be any suitable device made in accordance with well known T-1 communications architecture, and/or architecture adapted for compatibility with other network and telecommunications architectures, protocols and topologies  
30   including, but not limited to, T-3, DS-3, OC-3C, OC-12C, OC-192C, FDDI, SONET, SCSI, TCP/IP, HiPPI and ATM. In addition, the computers 102 and 110 may be networked together and with other computers through appropriate ethernet

cards or other suitable networking techniques. The respective manufacturer, Fibre Channel Consortium and I2O Special Interest Group reference design data sheets and materials describing the detailed operating capabilities and specifications of each of the foregoing components are hereby incorporated by reference in their entirety. Also, further information concerning possible subsystems and connection protocols for the server are described in the referenced Dellacona applications.

Figures 2A, 2B and 2C illustrate generally three types of system configurations which can be addressed in accordance with the architecture of the present invention. Each Figure is diagrammatically illustrative of a chassis 300 with a power supply section 302 and a diagnostics section 304. In addition, the chassis includes a processing section 306 such one or more of the servers 100 of Figure 1, and /or a storage section 308 such as one or more of the mass storage modules 202 of Figure 1. The size of the processing and storage sections can be readily configured for a particular application.

For example, Figure 2A illustrates an arrangement for a high volume server such as an application server, a movie server (e.g., for video streaming to multiple users) or communications in conjunction with a carrier-class switch. It can be seen that the processing section 306 includes ten slots, each preferably representing a single board computer. In contrast, the storage section 308 has only five storage slots, each representing a high speed, high capacity storage device such as a disk drive. Thus, the processing function is stressed over the storage function. On the other hand, Figure 2B illustrates an arrangement which might be suitable for a web-server or web-hosting. Here, the processing section 306 comprises two slots of the chassis whereas the storage section 308 comprises fifteen slots. Figure 2C illustrates a chassis arrangement suitable primarily for storage expansion. Here, there is no processing section and the chassis is devoted to storage.

Figure 3 illustrates an embodiment of the mass storage array 200 of Figure 1 in greater detail as it could be configured with a single computer server 100 and one or more mass storage modules 202. Referring to Figure 2, the mass storage array 200 includes one or more modules 202A, 202B ... 202n, each of which is preferably identical architecturally, although they may contain different numbers of storage devices and different types of storage devices. Each module 202 includes bypass



Each module 202 has a module or chassis bypass circuit board 216 in the communication path of the Fibre Channel controller 104. The chassis bypass circuit board 216 of module 202A is provided with an optical input/ output connector 218 for outputting electrical signals from the module 202A as light signals and for inputting light signals into the module 202A as electrical signals. Likewise, the modules 202B .. 202n have chassis bypass circuit boards 216 with associated input/output connectors 218 (not shown). As illustrated, the input/output connector 218 of the chassis bypass circuit board of module 202A is adapted to be connected via a light transmission medium such as optical fibers 220 to the optical input/output connector 218 (not shown) of the chassis bypass circuit board of module 202B.

As was previously noted, the controller 104 preferably is a conventional  
15 Fibre Channel Controller (FCC) which operates on a Fibre Channel protocol, and  
preferably is an arbitrated dual channel Fiber Channel Controller. The controller  
104 provides a dual channel communication path within each module 202 between  
the computer 102 and each of the operable storage devices 212. As is described  
hereinafter in greater detail, the bypass circuit boards 210 ensure that the  
20 communication path is complete even if a storage device 212 is inoperable (i.e., is  
not operable at or above some minimum level as is hereinafter described in greater  
detail) or has been removed from the connector on the bypass circuit board.

In that regard, each of the storage devices 212 is preferably a high speed, high capacity, conventionally available disk drive which is removably connected to its associated bypass circuit board 210. Preferably, the disk drive plugs into a connector on the bypass circuit board 210 so that it can be readily removed and replaced or so that drives may be added to empty bypass circuit board positions, as needed to expand the storage capacity of the module. Each bypass circuit board 210 includes circuits which connect the controller 104 to the disk drive 212 when the disk drive is plugged in and is conveying to the bypass circuit board that it is operable at a certain minimum level. On the other hand, the bypass circuit board 210 connects the controller 104 directly to the next bypass circuit board, bypassing

the disk drive 212, when the disk drive is not plugged in or is not operating at or above the minimum level. Any suitable, conventional disk drive of the type that runs self-diagnostics and provides a fault/no fault output signal may be used for this purpose.

5           With further reference to Figure 3, by way of example, the bypass circuit board 210A directs communications to the associated storage device 212A and then to the next bypass circuit board 210B when an operable storage device 212A is connected to the bypass circuit board 210A. When the storage device 212A is inoperable, i.e. not operating at some minimum satisfactory level, or when there is  
10   no storage device connected to the bypass circuit board 210A, e.g., if the storage device is removed for replacement, the dual channel communication path proceeds through the bypass circuit board 210A without interruption, i.e., bypasses the storage device connector. While one embodiment of a bypass circuit board to accomplish the foregoing connection and bypass functions is described hereinafter in greater  
15   detail, it will be appreciated by one skilled in the art that these functions can be accomplished in any suitable conventional manner by electronic switching circuits controlled by the fault/no fault signal from the storage device.

          With continued reference to Figure 3, the module or chassis bypass circuit board 216 provides functions similar to the storage or drive bypass circuit boards  
20   210 in the sense that they either route communications back to the controller directly if there is no additional module 202B connected to the module 202A or they route communications to the next module 202B if one is connected and signals are being received. In the case of a single module 202, therefore, there is a dual channel communications loop entirely within the module by which the computer 102  
25   communicates with each of the disk drives 212. On the other hand, by virtue of the chassis bypass circuit board 216, the dual channel communications loop extends to each of the disk drives 212 in the next module when one is connected. Specifically, if there is an additional module 202B connected to module 202A as illustrated, then communications from the computer 102 of module 202A are directed by the module  
30   bypass circuit board 216 to the next module 202B via connectors 218 and optical fibers 220 so that module 202B is within and traversed by the dual Fibre Channel communication loop.

0050257206400

It will be appreciated that with the above described architecture, individual mass storage device modules 202 of the mass storage array 200 may be expanded internally by adding disk drives or other suitable storage devices, and bad disk drives may be replaced without affecting the operation of the rest of the module or the system it is used in. This provides an extremely versatile hot swappable, hot expandable, mass storage device array. In addition, when the demands of the system exceed the capacity of a single module, an additional module may be added, again without interrupting the operation of the rest of the array or its system.

As was explained above, the Fibre Channel controller provides a dual path through each module of the mass storage array 200. In accordance with the present invention, the system can be configured so that the dual path is a single path which traverses the mass storage array twice or two independent paths as is illustrated in Figures 4A and 4B.

Referring now to Figures 4A and 4B, the Fibre Channel controller (FCC) 104 includes two output paths A and B. Each path traverses the mass storage module 202 as illustrated, communicating with each of the present and operable storage devices 212 and returning to the FCC. In the Figure 4A embodiment, the A path returning to the FCC is looped back to the B output path. Accordingly, a single continuous path traverses the mass storage module twice. In the Figure 4B embodiment, the A return path is not looped back to the B output path. Accordingly, two paths traverse the mass storage module and can be used independently. This latter embodiment provides a second path in the event that one path fails.

It will be appreciated by one skilled in the art that the Figure 4A and 4B embodiments provide redundancy and fault tolerance. The Figure 4A embodiment is somewhat simpler to implement because only one set of chips is necessary to provide the single Fibre Channel capability required for the single loop. Still, if one of the loops is broken or encounters some other fault, that loop can be bypassed within the controller, and the other loop is still available. Similarly, while the Figure 4B embodiment may be more complex and expensive to implement, it provides two independently addressable loops for fault tolerance and redundancy, but also provides significantly greater communications bandwidth.

Figure 5 illustrates a further configuration which is made possible by the system architecture of the present invention. In the Figure 5 embodiment, two servers are connected to a mass storage array, a first module of which is illustrated. The servers 100A and 100B are connected to Fibre Channel bypass boards 222A and 222B, respectively. The operation of the computers 102A and 102B may be sensed by the Fibre Channel bypass boards, for example by sensing signal flow as with the chassis bypass board, so that the computers 102A and 102B can operate together or, in the case of a fault, separately. As was previously mentioned, each server may have flash memory, and in the embodiment of Figure 5, each computer 102A and 102B has its own flash memory with its operating system stored therein. In this fashion, the computers can boot independently from its associated flash memory rather than from the shared memory or other disk arrangement.

It will be appreciated that the system configuration illustrated in Figure 5 provides two servers and thus increased processing power. In addition, one server provides backup to the other in the event of failure, thereby providing increased fault tolerance and reliability. In addition, the architecture of Figure 5 permits the two servers to communicate with each other at Fibre Channel speeds, higher than could be achieved with 100BaseT LAN, using IP protocol.

It can be seen that system architecture according to the present invention lends itself to a wide variety of configurations to accommodate a variety of applications. Figure 6 illustrates one configuration wherein two servers A and B are connected in a suitable local area network (LAN) configuration with access to the internet. Server A has its own mass storage array 200 as does server B. It will be appreciated that several servers and storage arrays can be networked in this fashion to provide a very powerful information or application server system particularly suitable for web server applications. In addition, this configuration, permits full duplication of computer/mass storage systems rather than sharing a mass storage system as with the embodiment of Figure 5. Moreover, while the Figure 6 embodiment illustrates communication over a LAN which typically may be a 100BaseT LAN, it will be appreciated that this may be a Fibre Channel LAN if rates in the gigabit range are desired. These features and advantages may be the ideal choice for critical processing applications such as web servers.

Figure 7 illustrates another configuration particularly useful for video streaming applications. In this embodiment, the diagnostics computer 110 also has indexing functions for controlling access to content servers 1, 2 and 3. The content servers 100 provide access to video stored in their associated storage arrays 200.

5 Alternatives, also particularly suitable for video streaming applications, are shown in Figures 8 and 9. In the Figure 8 embodiment duplicated index servers separate from the diagnostics computer are provided with their own storage arrays. In the illustrated embodiment there are two index servers 1 and 2 and there are three content servers 1, 2 and 3, each with an associated storage array. In the Figure 9  
10 embodiment, the functions of the index servers are distributed among the content servers so that there are multiple index/content servers 1-5. One skilled in the art will appreciate that the Figure 8 approach uses an architecture where the index server or servers coordinate content streaming from the content servers whereas in Figure 9 the index and content server functionality is distributed across all servers,  
15 providing extensive scalability.

Figures 10 and 11 are functional block diagrams which generally illustrate the storage and chassis bypass circuit boards 210 and 216, respectively, in greater detail. Referring to Figure 10, the storage bypass circuit board 210 includes a bypass board backplane connector 230 arranged to plug into a connector on the  
20 backplane generally indicated at 232. The A and B signal paths coming from a previous bypass circuit or directly from the Fibre Channel controller are connected through the backplane to the bypass board via the backplane connector 230. Similarly, the A and B signal paths emerge from the bypass board 210 via the bypass board backplane connector 230.

25 The A signal path from the backplane connector 230 is connected to a suitable conventional electronic switch 234. The B signal path from the backplane connector 230 likewise is connected to an electronic switch 236. The A and B signal paths from electronic switches 234 and 236 are connected to a bypass board storage card or drive connector 238 where they are routed to the storage device (e.g., a disk  
30 drive) 212.

The return A signal path from the bypass board drive connector 238 is connected to the switch 234, and the return B signal path from the connector 238 is

00502572.061200

5

10

20

25

30

the switches 234A and 236A to the transceiver 218 where they are converted to light signals and transmitted over the optical fibers forming the A and B paths to the next chassis. Similarly, light signals returned from the next chassis on the A and B paths are converted back to electrical signals by the transceiver 218 and returned along the A and B paths through the switches 234A and 236A to the connector 230A and onto the next bypass circuit or the Fibre Channel controller.

Typically, each of the modules 202 of a mass storage device array such as the one shown in Figure 3 is enclosed in a housing or, if mounted in a rack, may be surrounded by other structures and/or other circuit boards. The server components such as the computers 102 and 110, the touch screen display 112, the Fibre Channel controller, and other system components such as communication access cards, ethernet cards, LAN components and the like may also be mounted within the same housing or on the same rack. Heating may therefore be a problem, particularly where the storage devices used in the modules are disk drives driven by motors. Accordingly, the preferred embodiment of the present invention as illustrated in Figures 12A and 12B includes a physical structure and arrangement of the various bypass circuit boards and storage devices that accommodates the circulation of cooling air throughout the module without minimal obstruction.

In accordance with one aspect of this physical structure of the invention, each of the drive bypass circuit boards 210 is a relatively thin circuit board. The circuit board is, however, unlike typical circuit boards built to receive a disk drive or other mass storage device. In such conventional circuit boards, the connector or plug which receives the plug-in disk drive typically is positioned so that the disk drive extends perpendicular to the plane of the board. For example, it is usual to have the bypass circuits and/or the communications paths between them on a backplane or midplane circuit board which extends across the module in a fashion similar to a computer motherboard which extends across the computer chassis and has connectors to receive various plug-in cards or boards. That arrangement creates an obstruction which makes it more difficult to effectively cool heat producing storage devices such as disk drives.

As is illustrated in Figures 12A and 12B, the present invention does not use either a circuit board midplane or a backplane structure. Instead, the connector on

the drive bypass board which receives the disk drive (e.g., the connector 238 in Figure 10) is at the edge of the board so that the drive, when plugged in, extends parallel to the circuit board. In addition, the drive bypass circuit board itself is not plugged into a backplane circuit board. Rather, a connector is provided on the edge of the board, preferably opposite the disk drive connector as illustrated at 239 in Figure 10, and that connector plugs into an individual connector which is mounted on a frame or other structural member of the chassis and which is wired or otherwise connected to similar connectors for the other disk bypass circuit boards.

Figures 12A shows a side view of the disk drive 212 plugged into the drive bypass circuit board 210 via connector 238, with the drive bypass circuit board plugged into connector 240 suitably mounted on structural members 242 of the chassis or rack containing the mass storage module and/or server and its associated components. Figure 12 B is an end view illustrating several bypass circuit boards 21 plugged into the connectors 240 which are in turn connected by screws or other suitable means to the structural members 242. Since there is no backplane which would normally make the connections between adjacent components, electrical or light connections generally indicated at 246 are suitably provided between the connectors 240 to provide the communications required, as illustrated, for example, in Figure 3 and 5.

It can be seen that between each connector 240 there is a space 244 through which air can readily be drawn or forced by a fan or other air circulation means as is necessary. Even if the structural members are part of a housing that surrounds the components, screening or other suitable openings can be provided so that the areas 244 permit sufficient air flow. It can also be seen that because the disk drive is plugged into the bypass circuit board so that their planes are parallel and not perpendicular, there is no obstruction of air flow.

It will also be appreciated that this arrangement is particularly suited for field service of the unit and is readily upgradeable. Connectors can be readily replaced in the field without the need to change a complete backplane or midplane board, and in some instances repairs of this sort can be carried out with little or no down time. In addition, the illustrated connection arrangement permits expansion without the limitations encountered when using a backplane or midplane with a fixed number of



Similar physical arrangements may be used to connect computers to their associated components to create the desired server configuration. For example, as illustrated in Figure 13A, a single board computer of the type previously described in connection with the description of the server 100 may connect to the communications card 106 through a mini compactPCI (CPCI ) backplane. In this case, the single board computer SBC may incorporate the Fibre Channel controller FCC. Likewise, in Figure 13B, the SBC may connect to an input/output unit I/O at the left rear connector of the CPCI backplane. In this case, the I/O may include the FCC functions. Again, it will be appreciated that the foregoing advantages are achieved with this sort of simple backplane or midplane structure which does not extend across the cabinet or rack.

The above-described exemplary embodiments are intended to be illustrative in all respects, rather than restrictive, of the present invention. Thus the present invention is capable of many variations in detailed implementation that can be derived from the description contained herein by a person skilled in the art. All such variations and modifications are considered to be within the scope and spirit of the present invention as defined by the following claims.